

Fast Path Message Transfer Agent
Bradley Taylor

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present invention relates to sending email messages between servers, and particularly to a fast path message transfer agent for these email messages.

Description of the Related Art

[0002] Figure 1 illustrates a simplified email system 100 that can send or receive messages over the Internet 103. System 100 typically uses Simple Mail Transfer Protocol (SMTP) to send messages between servers 101 and 102. Clients 104-106 can use server 101 to route their email, whereas clients 107-109 can use server 102 to route their email.

[0003] Servers 101 and 102 temporarily store and re-route the email messages from clients 104-109 to the appropriate destinations. Specifically, Message Transfer Agents (MTAs) 110 and 111, installed in servers 101 and 102, respectively, can route messages according to addresses designated in the email. MTAs can use retry logic and queues (explained in further detail below) to efficiently direct the messages.

[0004] Figures 2A and 2B illustrate a flow chart of a standard embodiment of an MTA operation. In step 201, the MTA, in a first email server, receives a network connection from a second email server. The MTA can then receive bytes of a message over this connection in step 202. In step 203, the MTA writes those bytes into a dynamic (e.g. DRAM or SRAM) memory to quickly capture those bytes and then stores those bytes into a non-volatile storage device until delivery to their final destination(s) (i.e. the clients associated with the first email server).

[0005] If all bytes of the message have not been received, as determined in step 204, then the process returns to step 202 to receive additional bytes of the message over the connection. If all bytes of the message have been received, then the MTA responds to the server that the message has been successfully received in step 205. Note that if a system failure (e.g. a full condition or power outage) occurs during steps 201-203, then the MTA can respond to the server with an error message, wherein the server can re-establish network connection at a later point in time to resend the message. In this case, the MTA can delete any bytes of the message that were written to the memory or stored in the non-volatile storage device.

[0006] Assuming successful receipt, the stored bytes for a message in the non-volatile storage device are now in a queue of messages to be re-routed to their destinations. When a message is next in queue, the MTA retrieves the message from the non-volatile storage in step 206. The MTA attempts to send the message to each destination designated by the message in step 207. If the message was not successfully delivered to all destinations, as determined by step 208, then the MTA can identify the failed destinations in step 210 and then retry delivering the message to the failed destinations after some delay in step 211. In one embodiment, the message and its failed destinations can be returned to a queue in the non-volatile storage device, wherein the process returns to step 206. If the message was successfully delivered to all destinations, then the message is removed from the non-volatile storage device in step 209 and the delivery process for that message ends.

[0007] Continuously storing and accessing messages on the non-volatile storage device undesirably increases the time of email delivery. Therefore, a need arises for a method of decreasing email delivery time.

SUMMARY OF THE INVENTION

[0008] A method of providing a fast path message transfer agent (MTA) is provided. A typical implementation of the fast path MTA can increase performance by 3X the speed of a standard MTA. The method includes receiving bytes of a message over a network connection and determining whether the number of bytes exceeds a predetermined threshold. If the number of bytes is less than a predetermined threshold, then the message is written only to memory. However, if the number of bytes exceeds the predetermined threshold, then the message is written to memory and a non-volatile storage device. In one embodiment, some of the bytes (e.g. up to the predetermined threshold) are written to memory, wherein the remainder of the bytes are stored into the non-volatile storage device.

[0009] Writing the message to the memory and the non-volatile storage device can further include determining whether all bytes of the message have been received. If not, then additional bytes of the message can be received over the network connection. The additional bytes can be written into the non-volatile storage device.

[0010] The method can further including accessing the message, sending the message to each destination, and determining whether the message was received successfully by each destination. If the message was received successfully by each destination, then the message can be removed from the memory (or from the memory and the non-volatile storage device, as appropriate) and a successful receipt of the message can be indicated.

[0011] If the message was not received successfully by each destination, then all failed destinations can be identified and the message can be stored in the non-volatile storage device. However, a successful receipt of the message can still be

indicated. The failed destinations can be retried after a delay until the message is successfully received. At this point, the message can be removed from the non-volatile storage device.

[0012] A computer program product is also provided. The product can include a computer usable medium having a computer readable program code embodied therein for providing a fast path message transfer agent. The computer readable program code can comprise computer readable program code that receives bytes of a message over a network connection and computer readable program code that determines if the number of bytes exceeds a predetermined threshold. If the number of bytes is less than the threshold, then the message is written only to memory. However, if the number of bytes exceeds the threshold, then the message is written to memory and a non-volatile storage device. The product can further comprise computer readable program code that writes some of the bytes (for example, up to the predetermined threshold) to memory and computer readable program code that stores a remainder of the bytes in the non-volatile storage device.

[0013] Another embodiment of a method for providing a MTA is provided. This method can include receiving a network connection from an email server, receiving addresses of any recipients, and determining whether connections can be formed to the recipients. If so, then bytes of a message can be received and sent to the recipients. If not, then the connections can be retried for a predetermined number of times.

[0014] If the bytes are received by the recipients, then the MTA can respond to the server that the message was successfully received by the recipients. On the other hand, if not all the bytes are received by the recipients, then the MTA can respond to the server that message transfer was not successful. In the case that retrying exceeds the predetermined number of times, then the

MTA can respond to the server that connections to the recipients were not successful.

[0015] Yet another embodiment of a method of providing a fast path MTA is provided. This method can include receiving a network connection from an email server, receiving bytes of a message over the network connection, and determining whether the number of bytes exceeds a predetermined threshold. In this embodiment, if the number of bytes does not exceed a predetermined threshold not, then the message is written only to a memory. However, if the number of bytes exceeds a predetermined threshold, then the message is written only to non-volatile storage.

[0016] If all bytes of the message have not been received, then additional bytes of the message can be received. However, if the total number of bytes exceeds the predetermined threshold, then the total number of bytes are stored in the non-volatile storage device and any bytes of the message written to memory are erased.

BRIEF DESCRIPTION OF THE DRAWINGS

[0017] Figure 1 illustrates a simplified block diagram of an email system that can send or receive messages over the Internet.

[0018] Figures 2A and 2B illustrate a flow chart of one embodiment of a standard Message Transfer Agent (MTA) operation.

[0019] Figures 3A and 3B illustrate a flow chart of one embodiment of a fast path MTA operation.

[0020] Figure 4 illustrates an embodiment in which the fast path MTA operation can be disabled and the standard MTA operation can be enabled.

[0021] Figure 5 illustrates another embodiment of a fast path MTA that eliminates the need for disk storage.

[0022] Figure 6 illustrates another embodiment of a MTA that uses one of a fast path or a standard path.

DETAILED DESCRIPTION OF THE DRAWINGS

[0023] Figure 3A and 3B illustrate a flow chart of one embodiment of a fast path MTA operation. In this operation, storing a message in a non-volatile storage device is limited to circumstances where the number of bytes exceeds a predetermined threshold. The predetermined threshold is set so that a majority of the messages can be written only to memory. In this manner, the delivery time can be significantly reduced compared to the standard MTA operation (Figures 2A and 2B) that must include access time to the non-volatile storage.

[0024] In step 301, the MTA receives a network connection from the email server. In step 302, the MTA receives bytes of a message over that connection. If the total number of bytes (i.e. bytes written to memory in combination with those bytes just-received) does not exceed a predetermined threshold, as determined in step 303, then the received bytes are written to memory in step 304. If all bytes of the message have not been received, as determined in step 305, then the MTA returns to step 302 to receive additional bytes.

[0025] If the total number of bytes exceeds a predetermined threshold in step 303, then the number of bytes up to the threshold are written to memory and the remainder of bytes is stored in a non-volatile storage device in step 306. If all bytes of the message have not been received, as determined in step 307, then the MTA receives additional bytes of the message over the connection in step 308. These additional bytes of the message are stored only onto the non-volatile storage device in step 309. Once all bytes of the message have been received (steps 305/307), the bytes for the message in memory (or in the

memory and on the non-volatile storage device) can be re-routed to their destinations.

[0026] Thus, the MTA can follow one of two processes depending on the predetermined threshold. In one embodiment, the predetermined threshold is set so that a majority of the messages can be written only to memory. For example, the threshold can be set to 32k, although other thresholds (such as 16k or 64k) can also be used depending on the projected size of the files. Because the MTA stores onto and accesses the non-volatile storage device infrequently (and in preferred cases, not at all), this process is significantly faster than the prior art process. In fact, for typical implementations, this fast path MTA process can be three times as fast as the prior art MTA process.

[0027] To re-route the message to its destination(s), the MTA accesses the message from memory (or accesses a portion of the message from memory and retrieves the remainder of the message from the non-volatile storage device) in step 310. The MTA attempts to send the message to each destination designated by the message in step 311. If the message was successfully delivered to all destinations, as determined by step 312, then the message can be removed from memory (or memory and the non-volatile storage device) in step 313. The MTA can respond to the server that the message was successfully received in step 314 and the delivery process for that message ends.

[0028] If the message was not successfully delivered to all destinations, as determined by step 312, then the MTA can identify the failed destinations in step 315 and then store the complete message (and the failed destinations) onto a non-volatile storage in step 316. In one embodiment, the message and its failed destinations can be placed in a queue in the non-volatile storage device. Note that, in addition to system failure or memory problems at the destination, a failed delivery

can occur because of a condition determined at the time of delivery. For example, if the message is an 8-bit message and the destination only supports 7-bit, then a bit conversion must be done before the message can be successfully delivered to that destination. The bit conversion of the 8-bit message can be done by a tool activated by the MTA, wherein after conversion the 7-bit message can be put in the queue of the non-volatile storage device. Once the message is stored onto the non-volatile storage device, the MTA can respond to the server that the message has been successfully received in step 317. The MTA can retry sending the message to the failed destinations after a predetermined delay in step 318.

[0029] If the message is not successfully received by all destinations, as determined in step 319, then the MTA can repeat steps 318 and 319. Note that in one embodiment, after a predetermined of retries, the process proceeds to step 320. Once a message is successfully received by all destinations, then the message can be removed from the non-volatile storage device in step 320 and the delivery process for that message ends. Note that receiving another network connection from an email server can occur at any time. Thus, one set of steps 301-320 can be interleaved with one or more other sets of steps 301-320, as needed.

[0030] In one embodiment, the fast path MTA can be disabled, thereby activating the standard path MTA described in Figures 2A and 2B. For example, referring to Figure 4, after receiving the bytes of the message over the connection in step 302 (Figure 3A), step 401 can determine if any predetermined conditions are present in the message. If no predetermined conditions are present, then the fast path MTA process can be implemented (step 402) by proceeding to step 303 (also Figure 3A). If predetermined conditions are present, then the standard path MTA

process can be implemented (step 403) by proceeding to step 203 (Figure 2A).

[0031] Predetermined conditions can include, by way of example and not limitation, enabled filtering (e.g. anti-virus, anti-spam, and content-filtering), disabled Lightweight Directory Access Protocol (LDAP) relaying (wherein LDAP includes a set of protocols for accessing information directories including email addresses etc.), enabled encryption, enabled LDAP "sender masquerading" (feature that refers to the LDAP directory to replace the sender's name with another entry, such as the "official" email address for the sender), enabled Realtime Blackhole List (RBL) (includes a listing of IP addresses whose owners refuse to stop spam), enabled "received for header" (feature that causes the MTA to take a message that is addressed to several destinations and transform the message into several messages, each message addressed to one of the several destinations), and enabled "sender check" (feature that ensures the sender's domain exists in the Domain Name System (DNS), an Internet service that translates domain names into IP addresses). Note that although the predetermined conditions indicated above relate to the configuration of the server, other predetermined conditions could exist to trigger disabling of the fast path MTA.

[0032] In accordance with another embodiment of the invention, the need for disk storage is completely eliminated. For example, referring to Figure 5, the MTA receives a network connection from the email server in step 501. In step 502, the MTA receives addresses of the recipients. If connections can be formed to all recipients, as determined in step 503, then the MTA receives the bytes of the message in step 504. These bytes can be sent to the recipients in step 505. Note that steps 504 and 505 could be performed substantially simultaneously or step 505 could follow completion of step 504. If the recipients received all the bytes

of the message, as determined in step 506, then the MTA can respond to the server that the messages was successfully received in step 507. If the recipients did not receive all bits, then the MTA can respond to the server that the message transfer was not successful in step 508. Note that if connections cannot be formed to all recipients in step 503, then either the MTA can respond to the server that connections to all recipients were not successful in step 509 or, alternatively, the MTA can retry connecting (i.e. returning to step 503) a predetermined number of times before responding to the server.

[0033] In accordance with another embodiment of the invention, the need for disk storage is eliminated in the fast path MTA. For example, referring to Figure 6, the MTA receives a network connection from the email server in step 601. In step 602, the MTA receives bytes of a message over that connection. If the total number of bytes (i.e. bytes written to memory in combination with those bytes just-received) does not exceed a predetermined threshold, as determined in step 603, then the received bytes are written to memory in step 604. If all bytes of the message have not been received, as determined in step 605, then the MTA returns to step 602 to receive additional bytes.

[0034] If the total number of bytes exceeds a predetermined threshold in step 603, then the fast path MTA is bypassed. Specifically, the bytes comprising the message are stored in a non-volatile storage device and any bytes of the message written to memory are erased in step 607. If all bytes of the message have not been received, as determined in step 608, then the MTA receives additional bytes of the message over the connection in step 609. These additional bytes of the message are stored only onto the non-volatile storage device in step 610. Once all bytes of the message have been received, as determined in step 608, the MTA can respond to the server that the message has been

successfully received in step 205 (Figure 2A). Steps 206-213 (Figure 2B) can then be followed in accordance with a standard MTA. Finally, assuming the fast path MTA is still active and all bytes for the message have been received, as determined in step 605, then steps 310-320 (Figure 3B) can be followed.

[0035] In one exemplary computer implemented embodiment, the MTA may be run on a variety of computer platforms including, for example: a SPARC station 20 or higher using the Sun™, Solaris™, or SPARC 2.6/2.7 operating system with at least 128 MB of RAM. In another embodiment, a PC platform (such as a Dell PowerEdge™ 1500) running either RedHat Linux 7.1 or Microsoft Windows™ 2000 can be used.

[0036] Although illustrative embodiments of the invention have been described in detail herein with reference to the accompanying figures, it is to be understood that the invention is not limited to those precise embodiments. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed. As such, many modifications and variations will be apparent to practitioners skilled in this art. For example, although a non-volatile storage device is discussed herein, any type of stable storage device could be used. In one embodiment, the non-volatile storage device could be a disk. Note that the predetermined threshold could be set to any value, even a value that would eliminate the need for the non-volatile storage device. Accordingly, it is intended that the scope of the invention be defined by the following Claims and their equivalents.